
Contents

1	Introduction	3
2	Quick Start	4
2.1	Download and Initial Build	4
2.2	Running the First 2D Example	5
2.3	Viewing the First 2D Example	6
2.4	Modifying the First Example	7
3	Basic Problem Setup	9
3.1	Contents of .rea file	9
3.2	Contents of .usr file	11
3.3	Contents of SIZE file	11
3.4	Memory Requirements	12
4	Simulation Capabilities	13
4.1	Governing Equations	13
4.1.1	Perturbation Equations	16
4.2	Geometry	16
4.2.1	Convention and Restrictions	16
4.2.2	Moving Geometry	17
4.3	Boundary Conditions	18
4.3.1	Fluid Velocity	19
4.3.2	Temperature	21
4.3.3	Passive scalar	21
4.3.4	Internal Boundary Conditions	21
4.4	Initial Conditions	22

4.5	Material Properties and Forcing Functions	23
5	Timestepping	25
5.1	Theoretical Background	25
5.2	Generalization of BDF k /EXT k	26
5.3	Stability Considerations	27
5.4	Characteristics Method	29
5.5	The Operator-Integration Factor Scheme	31
5.6	Extension to Navier-Stokes	33

Chapter 1. Introduction

This manual provides an introduction to the use of Nek5000, which is a simulation code for analyzing unsteady incompressible fluid flow with thermal and passive scalar transport. Nek5000 can treat general two- and three-dimensional domains described by isoparametric quad or hex elements. In addition, it can be used to compute axisymmetric flows. It is a time-stepping based code and does not currently support steady-state solvers, other than steady Stokes flow and steady heat conduction.

Nek5000 is based on the Nekton 2.0 spectral element code written by Paul Fischer, Lee Ho, and Einar Rønquist in 1986-1991, with technical input from A. Patera and Y. Maday. Nekton 2.0 was the first three-dimensional spectral element code and one of the first commercially available codes for distributed-memory parallel processors. Nek5000 development was continued from the 90s through the present by P. Fischer, Jerry Kruse, Julie Mullen, Henry Tufo, James Lottes, Chaman Verma, and Stefan Kerkemeier.

Note: This is an early draft version of the manual, very much under construction. It has been posted in its current form to provide a users with a set of basic instructions to get started with Nek5000. A more complete, but slightly dated, starting point for using the code can be found in

www.mcs.anl.gov/~fischer/primer.pdf

Material from the primer is being incorporated into the current document, which will be continually updated online.

Chapter 2. Quick Start

This chapter provides a quick overview to using Nek5000 for some basic flow problems provided in the `.../examples` directory.

2.1 Download and Initial Build

Nek5000 runs under linux or any linux-like OS such as MAC, AIX, BG, Cray etc. The source is maintained in an svn repository and can be downloaded from the Nek5000 homepage (google nek5000) or, on linux systems, with the svn checkout command:

```
svn co https://svn.mcs.anl.gov/repos/nek5 nek5_svn
```

After downloading, build the tools by typing

```
cd nek5_svn/trunk/tools
maketools all
```

which will put the tools `genbox`, `genmap`, `n2to3`, `postx`, `prex`, and `pretex` in the top-level `/bin` directory (and will create `/bin` if it does not exist). It may be necessary to edit the `maketools` file to change the compilers (e.g., to `pgf77/pgcc` or `ifort/icc`). However, the default `gfortran/gcc` is generally fine. ¹

Chapter 3. Basic Problem Setup

Nek5000 consists of three principal modules: the preprocessor **prex**, the solver **nek5000**, and the postprocessor **postx**. **prex** and **postx** are based upon an X-windows GUI. Nek5000 output formats can also be read by the parallel visualization package VisIt developed by Hank Childs and colleagues at LLNL/LBNL. VisIt is mandatory for large problems (e.g., more than 100,000 spectral elements). Nek5000 is written in f77 and C. It uses MPI for message passing (but can be compiled without MPI for serial applications) and some LAPACK routines for eigenvalue computations (depending on the particular solver employed). In addition, it can be optionally coupled with MOAB, which provides an interface to meshes generated with CUBIT.

Each simulation is defined by three files, the `.rea` file, the `.usr` file, and the `SIZE` file. In addition, there is a derived `.map` file that is generated from the `.rea` file by running **genmap**. Suppose you are doing a calculation called “shear.” The key files defining the problem would be `shear.rea`, `shear.map`, and `shear usr`. `SIZE` controls (at compile time) the polynomial degree used in the simulation, as well as the space dimension $d = 2$ or 3 .

This chapter provides an introduction to the basic files required to set up a Nek5000 simulation.

3.1 Contents of `.rea` file

The `.rea` file consists of several sections:

parameters These control the runtime parameters such as viscosity, conductivity, number of steps, timestep size, order of the timestepping, frequency of output, iteration tolerances, flow rate, filter strength, etc. There are also a number of free parameters that the user can use as handles to be passed into the user defined routines in the `.usr` file.

logicals These determine whether one is computing a steady or unsteady solution, whether advection is turned on, etc.

geometry The geometry is specified in an arcane format specifying the xyz locations of each of the eight points for each element, or the xy locations of each of the four points for each element in 2D. (For several reasons, this format is due to be changed in the future.)

curvature This section describes the deformation for elements that are curved. Currently-supported curved side or edge definitions include “C” for circles, “s” for spheres, and “m” for midside-node positions associated with quadratic edge displacement.

boundary conditions Boundary conditions (BCs) are specified for each face of each element, for each **field** (velocity, temperature, passive scalar #1, etc.). A common BC is **P**, which indicates that an element face is connected to another element to establish a periodic BC. Many of the BCs support either a constant specification or a user defined specification which may be an arbitrary function. For example, a constant Dirichlet BC for velocity is specified by **V**, while a user defined BC is specified by **v**. This upper/lower-case distinction is used for all cases. There are about 70 different types of boundary conditions in all, including free-surface, moving boundary, heat flux, convective cooling, etc.

restart conditions Here, one can specify a file to use as an initial condition. The initial condition need not be of the same polynomial order as the current simulation. One can also specify that, for example, the velocity is to come from one file and the temperature from another. The initial time is taken from the last specified restart file, but this can be overridden.

output specifications Outputs are discussed in a separate section below.

It is important to note that Nek5000 currently supports two input file formats, ascii and binary. The **.rea** file format described above is ascii. For the binary format, all sections of the **.rea** file having storage requirements that scale with number of elements (i.e., geometry, curvature, and boundary conditions) are moved to a second, **.re2**, file and written in binary. The remaining sections continue to reside in the **.rea** file. The distinction between the ascii and binary formats is indicated in the **.rea** file by having a negative number of elements. There are converters, **reatore2** and **re2torea**, in the Nek5000 tools directory to change between formats. The binary file

format is the default and important for i/o performance when the number of elements is large (> 100000, say).

3.2 Contents of .usr file

The most important interface to Nek5000 is the set of fortran sub-routines that are contained in the `.usr` file. This file allows direct access to all runtime variables. Here, the user may specify spatially varying properties (e.g., viscosity), volumetric heating sources, body forces, and so forth. One can also specify arbitrary initial and boundary conditions through the routines `useric()` and `userbc()`. The routine `userchk()` allows the user to interrogate the solution at the end of each timestep for diagnostic purposes. The `.usr` files provided in the `/examples/...` directories illustrate several of the more common analysis tools. For instance, there are utilities for computing the time average of u , u^2 , etc. so that one can analyze mean and rms distributions with the postprocessor. There are routines for computing the vorticity or the scalar λ_2 for vortex identification, and so forth.

3.3 Contents of SIZE file

The SIZE file governs the memory allocation for most of the arrays in Nek5000, with the exception of those required by the C utilities. The primary parameters of interest in SIZE are:

ldim = 2 or 3. This must be set to 2 for two-dimensional or axisymmetric simulations (the latter only partially supported) or to 3 for three-dimensional simulations.

lx1 controls the polynomial order of the approximation, $N=lx1-1$.

lxd controls the polynomial order of the integration for convective terms. Generally, $lxd=3*lx1/2$. On some platforms, however, it is important for memory access performance that $lx1$ and lxd be even.

lx2 = $lx1$ or $lx1-2$. This determines the formulation for the Navier-Stokes solver (i.e., the choice between the \mathbb{P}_N - \mathbb{P}_N or \mathbb{P}_N - \mathbb{P}_{N-2} methods) and the approximation order for the pressure, $lx2-1$.

lelt determines the *maximum* number of elements *per processor*.

3.4 Memory Requirements

Per-processor memory requirements for Nek5000 scale roughly as 400 8-byte words per allocated gridpoint. The number of *allocated* gridpoints per processor is $n_{\max} = lx1 * ly1 * lz1 * lelt$. (For 3D, $lz1 = ly1 = lx1$; for 2D, $lz1 = 1$, $ly1 = lx1$.) If required for a particular simulation, more memory may be made available by using additional processors. For example, suppose one needed to run a simulation with 6000 elements of order $N = 9$. To leading order, the total memory requirements would be $\approx E(N+1)^3 points \times 400(wds/pt) \times 8bytes/wd = 6000 \times 10^3 \times 400 \times 8 = 19.2$ GB. Assuming there is 400 MB of memory per core available to the user (after accounting for OS requirements), then one could run this simulation with $P \geq 19,200MB/(400MB/proc) = 48$ processors. To do so, it would be necessary to set $lelt \geq 6000/48 = 125$.

We note two other parameters of interest in the parallel context:

lp, the maximum number of processors that can be used.

lelg, an upper bound on the number of elements in the simulation.

There is a slight memory penalty associated with these variables, so one generally does not want to have them excessively large. It is common, however, to have **lp** be as large as anticipated for a given case so that the executable can be run without recompiling on any admissible number of processors ($P_{\text{mem}} \leq P \leq E$, where P_{mem} is the value computed above).

Chapter 4. Simulation Capabilities

This chapter describes the capabilities of **Nek5000** by introducing the set of governing equations together with boundary conditions, initial conditions, and material property/forcing functions that the program can solve numerically.

4.1 Governing Equations

Nek5000 is designed to simulate laminar, transitional, and turbulent incompressible or low Mach-number flows with heat transfer and species transport. It is also suitable for incompressible magnetohydrodynamics (MHD), described in Section ??.

Nek5000 solves the unsteady incompressible two-dimensional, axisymmetric, or three-dimensional Stokes or Navier-Stokes equations with forced or natural convection heat transfer in both stationary (fixed) or time-dependent geometry. The solution variables are the fluid velocity $\mathbf{u} = (u_x, u_y, u_z)$, the pressure p , the temperature T , independent passive scalar fields ϕ_i , $i=1,2,\dots$, mesh velocity $\mathbf{w} = (w_x, w_y, w_z)$ (for time-dependent geometry), and magnetic field $\mathbf{B} = (B_x, B_y, B_z)$ (for MHD). All of the above field variables are functions of space $\mathbf{x} = (x, y, z)$ and time t in domains Ω_f and/or Ω_s defined in Fig. 4.1 The governing equations used are: the momentum equations in Ω_f ,

$$\rho(\partial_t \mathbf{u} + \mathbf{u} \cdot \nabla \mathbf{u}) = -\nabla p + \nabla \cdot [\mu(\nabla \mathbf{u} + (\nabla \mathbf{u})^T)] + \rho \mathbf{f} , \quad (4.1)$$

the continuity (mass conservation) equation in Ω_f ,

$$\nabla \cdot \mathbf{u} = 0 , \quad (4.2)$$

the energy equation in $\Omega_f \cup \Omega_s$,

$$\rho c_p (\partial_t T + \mathbf{u} \cdot \nabla T) = \nabla \cdot (k \nabla T) + q_{vol} , \quad (4.3)$$

and a convection-diffusion equation for each passive scalar ϕ_i , $i=1,2,\dots$ in $\Omega_f \cup \Omega_s$

$$(\rho c_p)_i (\partial_t \phi_i + \mathbf{u} \cdot \nabla \phi_i) = \nabla \cdot (k_i \nabla \phi_i) + (q_{vol})_i . \quad (4.4)$$

Figure 4.1: Computational domain showing respective fluid and solid subdomains, Ω_f and Ω_s .

The above governing equations are subject to boundary conditions and initial conditions described in the following sections.

It should be noted that Nek5000 can solve not only the above fully coupled system, but also various subsets of these equations. Namely, Nek5000 recognizes the following subsets.

1. Fluid flow only; i.e., the momentum and continuity equations only.
2. Fluid flow and heat transfer only; i.e., the momentum, continuity and energy equations only.
3. Fluid flow, heat transfer, and an arbitrary subset of passive scalar fields; i.e., the momentum, continuity, energy and convective-diffusive equations for ϕ_i
4. Fluid flow with the nonstress formulation; i.e., the momentum equation, in which the term $(\nabla \mathbf{u})^T$ in the viscous contribution is set to zero, the continuity equation, and with or without the energy and convective-diffusive equations.
5. Low Mach-number flows where density may vary due to thermal-dilation effects.
6. Unsteady Stokes flow; i.e., the momentum equation, in which the nonlinear term $\mathbf{u} \cdot \nabla \mathbf{u}$ is set to zero, and the continuity equation.
7. Steady Stokes flow; i.e., the momentum and continuity equations with the nonlinear term $\mathbf{u} \cdot \nabla \mathbf{u}$ and transient term $\partial_t \mathbf{u}$ set to zero.
8. Transient heat transfer; i.e., the energy equation only with prescribed steady or unsteady velocity.
9. Steady conduction heat transfer; i.e., only the energy equation in which the advection term $\mathbf{u} \cdot \nabla T$ and the transient term $\partial_t T$ are zero.

Any combination of these equation characteristics is permissible with the following restrictions. First, the equation must be set to unsteady if it is time-dependent or if there is any type of advection. For these cases, the steady-state (if it exists) is found as stable evolution of the initial-value-problem. Secondly, the stress formulation must be selected if the geometry is time-dependent. In addition,

stress formulation must be employed if there are traction boundary conditions applied on any fluid boundary, or if any mixed velocity/traction boundaries, such as symmetry and outflow/n, are not aligned with either one of the Cartesian x, y or z axes. Using the minimum set of equations possible for a given application clearly saves computational time as Nek5000 solves only those equations specified.

4.1.1 Linearized Equations

In addition to the basic evolution equations described above, Nek5000 provides support for the evolution of small perturbations about a base state by solving the *linearized equations*

$$\rho(\partial_t \mathbf{u}'_i + \mathbf{u} \cdot \nabla \mathbf{u}'_i + \mathbf{u}'_i \cdot \nabla \mathbf{u}) = -\nabla p'_i + \mu \nabla^2 \mathbf{u}'_i, \quad \nabla \cdot \mathbf{u}'_i = 0, \quad (4.5)$$

for multiple perturbation fields $i = 1, 2, \dots$ subject to different initial conditions and (typically) homogeneous boundary conditions. These solutions can be evolved concurrently with the base fields (\mathbf{u}, p, T) . There is also support for computing perturbation solutions to the Boussinesq equations for natural convection. Calculations such as these can be used to estimate Lyapunov exponents of chaotic flows, etc.

4.2 Geometry

4.2.1 Convention and Restrictions

The domain in which the fluid flow/heat transfer problem is solved consists of two distinct subdomains. The first subdomain is that part of the region occupied by fluid, denoted Ω_f , while the second subdomain is that part of the region occupied by a solid, denoted Ω_s . These two subdomains are depicted in Fig. 2.1 for several different geometries. The entire domain is denoted as $D = \Omega_f \cup \Omega_s$. The fluid problem is solved in the domain Ω_f , while the temperature in the energy equation is solved in the entire domain; the passive scalars can be solved in either the fluid or the entire domain.

We denote the entire boundary of Ω_f as $\partial\Omega_f$, that part of the boundary of Ω_f which is not shared by Ω_s as $\partial\Omega'_f$, and that part of the boundary of Ω_f which is shared by Ω_s as $\partial\overline{\Omega}_f$ (note $\partial\Omega_f = \partial\Omega'_f \cup \partial\overline{\Omega}_f$). In addition, $\partial\Omega_s, \partial\Omega'_s$ and $\partial\overline{\Omega}_s$ are analogously defined. These distinct portions of the domain boundary are illustrated in Fig. 2.1. The restrictions on the domain for Nek5000 are itemized below.

- The domain $\Omega = \Omega_f \cup \Omega_s$ must correspond either to a planar (Cartesian) two-dimensional geometry, or to the cross-section of an axisymmetric region specified by revolution of the cross-section about a specified axis, or by a (Cartesian) three-dimensional geometry.
- For two-dimensional and axisymmetric geometries, the boundaries of both subdomains, $\partial\Omega_f$ and $\partial\Omega_s$, must be representable as (or at least approximated by) the union of straight line segments, splines, or circular arcs.
- Nek5000 can interpret a two-dimensional image as either a planar Cartesian geometry, or the cross-section of an axisymmetric body. In the case of the latter, it is assumed that the y -direction is the radial direction, that is, the axis of revolution is at $y=0$. Although an axisymmetric geometry is, in fact, three-dimensional, Nek5000 can assume that the field variables are also axisymmetric (that is, do not depend on azimuth, but only y , that is, radius, x , and t), thus reducing the relevant equations to "two-dimensional" form.
- In the geometry generation using PRENEK, a three-dimensional geometry must correspond to a series of layers of planar two-dimensional geometries. Each of those layers must adhere to the rules governing planar two-dimensional geometries. In addition, it is assumed that the parallel planes which separate the layers of elements are oriented parallel to the x - y plane, and that the first such plane is located at $z = 0$. Subsequent planes are located at arbitrarily increasing values of z . Note that this is a PRENEK restriction as Nek5000 can readily accept general three-dimensional geometries.

Fully general three-dimensional meshes generated by other softwares packages can be input to PRENEK as import meshes. Currently, meshes generated by PRE-BFC and I-DEAS can be accepted by PRENEK.

4.2.2 Moving Geometry

If the imposed boundary conditions allow for motion of the boundary during the solution period (for example, moving walls, free-surfaces, melting fronts, fluid layers), then the geometry of the computational domain is automatically considered in Nek5000 as being time-dependent.

For time-dependent geometry problems, a mesh velocity \mathbf{w} is defined at each collocation point of the computational domain (mesh) to characterize the deformation of the mesh. In the solution of the mesh velocity, the value of the mesh velocity at the moving boundaries are first computed using appropriate kinematic conditions (for free-surfaces, moving walls and fluid layers) or dynamic conditions (for melting fronts). On all other external boundaries, the normal mesh velocity on the boundary is always set to zero. In the tangential direction, either a zero tangential velocity condition or a zero tangential traction condition is imposed; this selection is automatically performed by Nek5000 based on the fluid and/or thermal boundary conditions specified on the boundary. However, under special circumstances the user may want to override the defaults set by Nek5000, this is described in the PRENEK manual in Section 5.7. If the zero tangential mesh velocity is imposed, then the mesh is fixed in space; if the zero traction condition is imposed, then the mesh can slide along the tangential directions on the boundary. The resulting boundary-value-problem for the mesh velocity is solved in Nek5000 using an elastostatic solver, with the Poisson ratio typically set to zero. The new mesh geometry is then computed by integrating the mesh velocity explicitly in time and updating the nodal coordinates of the collocation points.

Note that the number of macro-elements, the order of the macro-elements and the topology of the mesh remain *unchanged* even though the geometry is time-dependent. The use of an arbitrary-Lagrangian-Eulerian description in Nek5000 ensures that the moving fronts are tracked with the minimum amount of mesh distortion; in addition, the elastostatic mesh solver can handle moderately large mesh distortion. However, it is the responsibility of the user to decide when a mesh would become "too deformed" and thus requires remeshing. The execution of the program will terminate when the mesh becomes unacceptable, that is, a one-to-one mapping between the physical coordinates and the isoparametric local coordinates for any macro-element no longer exists.

4.3 Boundary Conditions

The boundary conditions for the governing equations given in the previous section are described in the following. Note that if the boundary conditions (for any field variable) are nonzero, the inhomogeneities can either be defined as constant, or as fortran functions of appropriate parameters such as space, time, temperature, etc. In this case, the user is responsible for using relevant variables in all

fortran function definitions.

4.3.1 Fluid Velocity

Two types of boundary conditions are applicable to the fluid velocity : essential (Dirichlet) boundary condition in which the velocity is specified; natural (Neumann) boundary condition in which the traction is specified. For segments that constitute the boundary $\partial\Omega_f$ (refer to Fig. 2.1), one of these two types of boundary conditions must be assigned to each component of the fluid velocity. The fluid boundary condition can be *all Dirichlet* if all velocity components of \mathbf{u} are specified; or it can be *all Neumann* if all traction components $\mathbf{t} = [-p\mathbf{I} + \mu(\nabla\mathbf{u} + (\nabla\mathbf{u})^T)] \cdot \mathbf{n}$, where \mathbf{I} is the identity tensor, \mathbf{n} is the unit normal and μ is the dynamic viscosity, are specified; or it can be *mixed Dirichlet/Neumann* if Dirichlet and Neumann conditions are selected for different velocity components. Examples for all Dirichlet, all Neumann and mixed Dirichlet/Neumann boundaries are wall, free-surface and symmetry, respectively. If the nonstress formulation is selected, then traction is not defined on the boundary. In this case, any Neumann boundary condition imposed must be homogeneous; i.e., equal to zero. In addition, mixed Dirichlet/Neumann boundaries must be aligned with one of the Cartesian axes.

For flow geometry which consists of a periodic repetition of a particular geometric unit, the periodic boundary conditions can be imposed, as illustrated in Fig. 2.2. For a fully-developed flow in such a configuration, one can effect great computational efficiencies by considering the problem in a single geometric unit (here taken to be of length L), and requiring periodicity of the field variables. Nek5000 requires that the pairs of sides (or faces, in the case of a three-dimensional mesh) identified as periodic be identical (i.e., that the geometry be periodic).

For an axisymmetric flow geometry, the axis boundary condition is provided for boundary segments that lie entirely on the axis of symmetry. This is essentially a symmetry (mixed Dirichlet/Neumann) boundary condition in which the normal velocity and the tangential traction are set to zero.

For free-surface boundary segments, the inhomogeneous traction boundary conditions involve both the surface tension coefficient σ and the mean curvature of the free surface.

4.3.2 Temperature

The three types of boundary conditions applicable to the temperature are: essential (Dirichlet) boundary condition in which the temperature is specified; natural (Neumann) boundary condition in which the heat flux is specified; and mixed (Robin) boundary condition in which the heat flux is dependent on the temperature on the boundary. For segments that constitute the boundary $\partial\Omega'_f \cup \partial\Omega'_s$ (refer to Fig. 2.1), one of the above three types of boundary conditions must be assigned to the temperature.

The two types of Robin boundary condition for temperature are : convection boundary conditions for which the heat flux into the domain depends on the heat transfer coefficient h_c and the difference between the environmental temperature T_∞ and the surface temperature; and radiation boundary conditions for which the heat flux into the domain depends on the Stefan-Boltzmann constant/view-factor product h_{rad} and the difference between the fourth power of the environmental temperature T_∞ and the fourth power of the surface temperature.

4.3.3 Passive scalar

The boundary conditions for the passive scalar fields are analogous to those used for the temperature field. Thus, the temperature boundary condition menu will reappear for each passive scalar field so that the user can specify an independent set of boundary conditions for each passive scalar field. The terminology and restrictions of the temperature equations are retained for the passive scalars, so that it is the responsibility of the user to convert the notation of the passive scalar parameters to their thermal analogues. For example, in the context of mass transfer, the user should recognize that the values specified for temperature and heat flux will represent concentration and mass flux, respectively.

4.3.4 Internal Boundary Conditions

In the spatial discretization, the entire computational domain is subdivided into macro-elements, the boundary segments shared by any two of these macro-elements in Ω_f and Ω_s are denoted as internal boundaries. For fluid flow analysis with a single-fluid system or heat transfer analysis without change-of-phase, internal boundary conditions are irrelevant as the corresponding field variables on these segments are part of the solution. However, for a multi-fluid system

and for heat transfer analysis with change-of-phase, special conditions are required at particular internal boundaries, as described in the following.

For a fluid system composed of multiple immiscible fluids, the boundary (and hence the identity) of each fluid must be tracked, and a jump in the normal traction exists at the fluid-fluid interface if the surface tension coefficient is nonzero. For this purpose, the interface between any two fluids of different identity must be defined as a special type of internal boundary, namely, a fluid layer; and the associated surface tension coefficient also needs to be specified.

In a heat transfer analysis with change-of-phase, Nek5000 assumes that both phases exist at the start of the solution, and that all solid-liquid interfaces are specified as special internal boundaries, namely, the melting fronts. If the fluid flow problem is considered, i.e., the energy equation is solved in conjunction with the momentum and continuity equations, then only the common boundary between the fluid and the solid (i.e., all or portion of $\partial\bar{\Omega}'_f$ in Fig. 2.1) can be defined as the melting front. In this case, segments on $\partial\bar{\Omega}'_f$ that belong to the dynamic melting/freezing interface need to be specified by the user. Nek5000 always assumes that the density of the two phases are the same (i.e., no Stefan flow); therefore at the melting front, the boundary condition for the fluid velocity is the same as that for a stationary wall, that is, all velocity components are zero. If no fluid flow is considered, i.e., only the energy equation is solved, then any internal boundary can be defined as a melting front. The temperature boundary condition at the melting front corresponds to a Dirichlet condition; that is, the entire segment maintains a constant temperature equal to the user-specified melting temperature T_{melt} throughout the solution. In addition, the volumetric latent heat of fusion ρL for the two phases, which is also assumed to be constant, should be specified.

4.4 Initial Conditions

For time-dependent problems Nek5000 allows the user to choose among the following types of initial conditions for the velocity, temperature and passive scalars:

- Zero initial conditions: default; if nothing is specified.
- Fortran function: This option allows the user to specify the initial condition as a fortran function, e.g., as a function of x ,

y and z .

- Presolv: For a temperature problem the presolv option gives the steady conduction solution as initial condition for the temperature. For a fluid problem this option *can* give the steady Stokes solution as the initial condition for the velocity *provided* that the classical splitting scheme is *not* used.
- Restart: this option allows the user to read in results from an earlier simulation, and use these as initial conditions.

A tabulated summary of the compatibility of these initial condition options with various other solution strategies/parameters is given in the appendix.

4.5 Material Properties and Forcing Functions

The following restrictions are imposed on the material property and forcing function parameters in the governing equations and initial and boundary conditions described in the above sections (here and in what follows time-independent implies no variation in time, whereas constant refers to uniform in space).

- ρ , the density, is taken to be time-independent and constant; however, in a multi-fluid system different fluids can have different value of constant density.
- μ , the dynamic viscosity can vary arbitrarily in time and space; it can also be a function of temperature (if the energy equation is included) and strainrate invariants (if the stress formulation is selected).
- σ , the surface-tension coefficient can vary arbitrarily in time and space; it can also be a function of temperature and passive scalars.
- $\bar{\beta}$, the effective thermal expansion coefficient, is assumed time-independent and constant.
- $\mathbf{f}(t)$, the body force per unit mass term can vary with time, space, temperature and passive scalars.
- ρc_p , the volumetric specific heat, can vary arbitrarily with time, space and temperature.

- ρL , the volumetric latent heat of fusion at a front, is taken to be time-independent and constant; however, different constants can be assigned to different fronts.
- k , the thermal conductivity, can vary with time, space and temperature.
- q_{vol} , the volumetric heat generation, can vary with time, space and temperature.
- h_c , the convection heat transfer coefficient, can vary with time, space and temperature.
- h_{rad} , the Stefan-Boltzmann constant/view-factor product, can vary with time, space and temperature.
- T_∞ , the environmental temperature, can vary with time and space.
- T_{melt} , the melting temperature at a front, is taken with time and space; however, different melting temperature can be assigned to different fronts.

In the solution of the governing equations together with the boundary and initial conditions, Nek5000 treats the above parameters as pure numerical values; their physical significance depends on the user's choice of units. The system of units used is arbitrary (MKS, English, CGS, etc.). However, the system chosen must be used consistently throughout. For instance, if the equations and geometry have been non-dimensionalized, the μ/ρ in the fluid momentum equation is in fact the inverse Reynolds number, whereas if the equations are dimensional, μ/ρ represents the kinematic viscosity with dimensions of *length*²/*time*.

Chapter 5. Timestepping

In this chapter, we consider theoretical and practical issues relating to the timestepping schemes used in Nek5000. In particular, we contrast the conditionally-stable BDF k /EXT k scheme with the “unconditionally-stable” characteristics-based timestepper.

5.1 Theoretical Background

We consider semi-implicit numerical solution of convection-dominated problems such as the incompressible Navier-Stokes equations

$$\begin{aligned}\frac{\partial \mathbf{u}}{\partial t} + \mathbf{u} \cdot \nabla \mathbf{u} &= -\nabla p + \frac{1}{Re} \nabla^2 \mathbf{u} \quad \text{in } \Omega, \\ \nabla \cdot \mathbf{u} &= 0 \quad \text{in } \Omega,\end{aligned}\tag{5.1}$$

and the convection-diffusion equation

$$\frac{\partial \phi}{\partial t} + \mathbf{c} \cdot \nabla \phi = \frac{1}{Pe} \nabla^2 \phi + f, \quad \text{in } \Omega.\tag{5.2}$$

subject to appropriate initial and boundary conditions on \mathbf{u} and ϕ . As with the Navier-Stokes case, we assume a divergence-free convecting velocity field, $\nabla \cdot \mathbf{c} = 0$.

We would like to construct a high-order scheme in time that treats the nonsymmetric convection term explicitly while treating the diffusion term implicitly. One possible approach is to use Crank-Nicolson (CN) for the diffusion term and third-order Adams-Bashforth (AB3) for the convection term. This approach amounts to computing ϕ^n by approximately integrating the time-derivative of ϕ over the interval $[t^{n-1}, t^n]$. The diffusion term is integrated using the trapezoidal rule, and the convection term is integrated by extrapolating the convection contribution over the interval $[t^{n-1}, t^n]$ based on available values from prior timesteps, t^{n-q} , $q = 1, \dots, 3$. The CN/AB3 scheme is

$$\begin{aligned}\frac{\phi^n - \phi^{n-1}}{\Delta t} &= \frac{1}{2} \frac{1}{Pe} (\nabla^2 \phi^n + \nabla^2 \phi^{n-1}) \\ &+ \frac{1}{12} (23 \mathbf{c} \cdot \nabla \phi|_{t^{n-1}} - 16 \mathbf{c} \cdot \nabla \phi|_{t^{n-2}} + 5 \mathbf{c} \cdot \nabla \phi|_{t^{n-3}}) + \frac{1}{2} (f^n + f^{n-1})\end{aligned}$$

If f happens to be dependent on ϕ^n , it can alternatively be lumped with the AB3 terms. The scheme is globally second-order accurate in time. We choose AB3 over AB2 because the work is essentially the same and because the stability region for AB3 encloses a significant portion of the imaginary axis and thus contains the eigenvalues of the convection-diffusion operator in the zero-diffusion limit.

An alternative approach to the CN-AB3 scheme is to use backward differentiation of order k (BDF k). Here the time-derivative of ϕ is approximated to order k , using a one-sided difference formula, and the convection and diffusion terms are evaluated at time t^n . To avoid implicit treatment of the nonsymmetric convection term, we approximate $\mathbf{c} \cdot \nabla \phi$ at t^n by extrapolation of order k (EXT k), using values from t^{n-q} , $q = 1, \dots, k$. The values for f can be extrapolated, if they are dependent on ϕ^n , or evaluated directly at t^n , if known. A BDF2/EXT2 formulation is given by

$$\frac{3\phi^n - 4\phi^{n-1} + \phi^{n-2}}{2\Delta t} = \frac{1}{Pe} \nabla^2 \phi^n - (2 \mathbf{c} \cdot \nabla \phi|_{t^{n-1}} - \mathbf{c} \cdot \nabla \phi|_{t^{n-2}}) + f^n \quad (5.4)$$

The scheme is globally second-order accurate in time. It is readily extended to k th-order accuracy by changing the second-order approximation to the time derivative on the left-hand side of (5.4) and using k th-order extrapolation for the convection term on the right. As with the CN/AB3 scheme, the cost for this approach is dominated by the work associated with the implicit solve. The added cost of increasing k is thus essentially nil. Note that it's not appropriate to use AB k coefficients on the right-hand side of (5.4) because (5.4) is an approximation to the *value* of the function at t^n , rather than an approximation to the *integral* of the function over the interval $[t^{n-1}, t^n]$.

5.2 Generalization of BDF k /EXT k

To extend the the BDF k /EXT k scheme to variable timestep sizes, we first rewrite (5.4) in a more general form:

$$\sum_{q=0}^k \beta_q \phi^{n-q} = - \sum_{p=1}^k \gamma_p \mathbf{c} \cdot \nabla \phi|_{t^{n-p}} + \mathcal{I}^n, \quad (5.5)$$

where \mathcal{I}^n accounts for all the implicit contributions. The terms β_q are the differentiation coefficients associated with the approximation of a derivative at time t^n based on known values at t^{n-q} , $q = 0, \dots, k$. The terms γ_p are interpolation coefficients required to approximate a function at time t^n based on known values at

t^{n-p} , $p = 1, \dots, k$. These coefficients are readily determined from any standard approach to polynomial interpolation at knot points $[t^{n-k}, \dots, t^{n-j}]$ ($j = 0$ or 1). A convenient code is the general-purpose routine for interpolation/differentiation weights, `FD_WEIGHTS`, developed by Bengt Fornberg [?]. Note that, because of potential round-off errors, it is probably better for long time integrations to use just the *change* in time values, $\delta^{n-q} := t^{n-q} - t^n$, rather than the time values themselves, to determine the interpolation/differentiation coefficients. Of course, the values of δ^{n-q} should be computed by summing Δt^{n-j} , $j = 0, \dots, q$, rather than taking the difference $t^{n-q} - t^n$.

5.3 Stability Considerations

The explicit treatment of the convection term leads to timestep size restrictions in the CN/AB3 and BDF k /EXT k schemes given in (5.3)–(5.5) above. Assuming a closed domain ($\mathbf{c} \cdot \mathbf{n} \equiv 0$), the linear operator will be skew-symmetric in the zero-diffusion limit. Assuming that the skew symmetry of the continuous (in space) operator is preserved in the numerical discretization, we should seek a temporal discretization that encloses a portion of the imaginary axis near the origin. Both AB3 and BDF3/EXT3 accomplish this.

Consider the linear convection problem

$$\frac{\partial \phi}{\partial t} + \mathbf{c} \cdot \nabla \phi = 0, \quad (5.6)$$

subject to appropriate initial and boundary conditions, and assume that it has been discretized in space to yield

$$\frac{d\phi}{dt} = -C\phi, \quad (5.7)$$

where ϕ is the vector of unknown basis coefficients and C is the discrete skew-symmetric convection operator. Discretizing (5.7) in time using AB3 yields

$$\phi^n = \phi^{n-1} - \frac{\Delta t}{12}(23C\phi^{n-1} - 16C\phi^{n-2} + 5C\phi^{n-3}), \quad (5.8)$$

where we have assumed that the timestep size, Δt , is constant and that the velocity field \mathbf{c} and, hence, C , is time invariant. The stability region for AB3 is found by considering the model problem $u_t = \lambda u$, inserting this into (5.8) ($\phi \leftarrow u$, $-C \leftarrow \lambda$), and solving for the locus of points in the complex ($\lambda \Delta t$)-plane such that the magnitude of u is unchanging (see, e.g., [?] and Appendix A.) Figure 5.1

Figure 5.1: Stability regions for (left) AB2 and BDF2/EXT2, (center) AB3 and BDF3/EXT3, and (right) AB3 and BDF2/EXT2a.

shows the stability region for AB3 and for BDF3/EXT3. We see that both schemes enclose a substantial portion of the imaginary axis near the origin. It is also possible to improve the stability of BDF2/EXT2 by using a three-term treatment of the second-order extrapolation. For example, for fixed Δt , the choice $\gamma_1 = 8/3$, $\gamma_2 = -7/3$, and $\gamma_3 = 2/3$, will provide a stability region that is similar to AB3, as seen in the right panel of Fig. 5.1.

The implications of Fig. 5.1 are that, for stability, one must choose $\Delta t \leq z_{\text{crit}} / \max |\lambda|$, where $z_{\text{crit}} = .7236$ for AB3 and $.6339$ for BDF3/EXT3. The value of $\max |\lambda|$ depends on the *spatial* discretization and on the computational mesh size. For classical centered finite differences (FD), one has

$$\max |\lambda| = \frac{|\mathbf{c}|}{\Delta x},$$

which leads to the well known Courant-Friedrichs-Lewy (CFL) number

$$CFL := \frac{|\mathbf{c}| \Delta t}{\Delta x} = \max |\lambda| \Delta t.$$

Thus, for FD + AB3, we must have $CFL \leq .7236$. For a spatial discretizations based on one-dimensional Fourier methods, the maximum eigenvalue is

$$\max |\lambda| = \pi \frac{|\mathbf{c}|}{\Delta x}$$

[?], which implies that $CFL \leq .7236/\pi$ is required for stability. For

spectral element methods, the bound is

$$\max |\lambda| = S \frac{|\mathbf{c}|}{\Delta x_{\min}},$$

where S is a value ranging between 1.52 and 1.16 as the polynomial order N is increased from 3 to ∞ [?]. Note that, for the spectral element method, Δx_{\min} scales as $O(N^{-2})$ as a result of the clustering of the Gauss-Lobatto-Legendre points near the ends of the reference interval. Schemes that avoid the CFL restriction are thus of greater interest in the SEM case.

5.4 Characteristics Method

The characteristics scheme due to Pironneau [?] avoids the CFL constraint in the semi-implicit formulation by rewriting the convective term in (5.2) as a material derivative to yield

$$\frac{D\phi}{Dt} = \frac{1}{Pe} \nabla^2 \phi + f, \quad (5.9)$$

where

$$\frac{D\phi}{Dt} := \frac{\partial \phi}{\partial t} + \mathbf{c} \cdot \nabla \phi \quad (5.10)$$

represents the change of ϕ along the path line (or characteristic) associated with the convecting velocity field \mathbf{c} . The basic idea behind this method is to apply BDF k to (5.9), to yield (e.g., for $k=2$):

$$\frac{3\phi^n - 4\tilde{\phi}^{n-1} + \tilde{\phi}^{n-2}}{2\Delta t} = \frac{1}{Pe} \nabla^2 \phi^n + f^n. \quad (5.11)$$

(As before, if f^n is dependent upon ϕ^n , it can be approximated using EXT k .) In (5.11), $\tilde{\phi}^{n-q}$ represents the value of ϕ at an earlier point in time (t^{n-q}) and at the point in space, \mathbf{X}^{n-q} , from which it originates prior to being convected forward by the velocity field. That is,

$$\tilde{\phi}^{n-q}(\mathbf{x}) := \phi(\mathbf{X}^{n-q}(\mathbf{x}), t^{n-q}). \quad (5.12)$$

Note that the domain of $\tilde{\phi}^{n-q}$ is coincident with $\Omega(t^n)$, whereas the domain of $\phi(\cdot, t^{n-q})$ must be coincident with $\Omega(t^{n-q})$. This is a subtle point, but the implications are that spatial discretization of (5.11) can proceed with the usual weighted residual formalism, that is, by multiplying (5.11) by a test function v , integrating over Ω , and restricting the search (trial) and test spaces to finite dimensional subspaces of $H_0^1(\Omega)$.

$$\begin{array}{ccc}
\begin{array}{c} \tilde{\phi}_j^{n-2} \\ \underline{\phi}_j \\ \mathbf{X}_j^{n-2} \end{array} & & \begin{array}{c} \tilde{\phi}_j^{n-1} \\ \underline{\phi}_j \\ \mathbf{X}_j^{n-1} \end{array} \\
& & \begin{array}{c} \phi_j^n \\ \mathbf{X}_j^n := \mathbf{x}_j \end{array}
\end{array}$$

Figure 5.2: BDF2 approximation of material derivative $D\phi/Dt$ along the characteristic emanating from \mathbf{x}_j . \mathbf{X}_j^{n-q} is the foot of the characteristic corresponding to $t^n - q\Delta t$.

To illustrate the principal implementation aspects of the characteristics scheme, we consider the fully discretized system (in time and space), assuming a nodal basis for ϕ and introducing gridpoints \mathbf{x}_j . Assume that the spatial discretization of (5.11) leads to the following linear system

$$\left(-\frac{1}{Pe}A + \frac{3}{2\Delta t}B\right)\underline{\phi}^n = B\underline{g}^n, \quad (5.13)$$

where A is the stiffness matrix, B is the mass matrix, and the right-hand side

$$\underline{g}^n := B\left(\underline{f}^n + \frac{4}{2\Delta t}\tilde{\phi}^{n-1} - \frac{1}{2\Delta t}\tilde{\phi}^{n-2}\right), \quad (5.14)$$

accounts for the data and the values of ϕ at the feet of the characteristics.

The central question with the characteristics scheme is how to find the values of ϕ at the foot of the characteristic associated with gridpoint \mathbf{x}_j , that is,

$$\tilde{\phi}_j^{n-q} := \phi(\mathbf{X}_j^{n-q}, t^{n-q}),$$

where \mathbf{X}_j^{n-q} is the foot of the characteristic emanating from \mathbf{x}_j , as illustrated in Fig. 5.2. The standard semi-Lagrangian formulation, suggested by Pironneau and followed by others, is,

- starting at each point \mathbf{x}_j , march backwards along the velocity characteristic for an amount of time $q\Delta t$,
- evaluate \mathbf{X}_j^{n-q} (the endpoint of the trajectory in the preceding step)
- then evaluate $\phi(\mathbf{X}^{n-q}, t^{n-q}) := \phi^{n-q}(\mathbf{X}^{n-q})$.

This obviously requires that one interpolate $\phi^{n-q}(\mathbf{x})$ at points that are not coincident with the gridpoints since, in general, one will not have $\mathbf{X}_j^{n-q} = \mathbf{x}_{j'}$ for any j' . In fact, the situation is somewhat worse. In addition to interpolating ϕ , it is necessary to interpolate the vector field \mathbf{c} at each point along the characteristic, in order to accurately march back along the trajectory. In low-order unstructured methods, most of the interpolation cost arises from determination of the cell containing the point of interest. In high-order methods, one has an additional cost because off-grid interpolation for an N th-order discretization in d space dimensions has a complexity of $O(N^d)$ *per interpolated value*. For a spectral element discretization comprising E elements of order N , there are $\sim EN^d$ gridpoints, and the semi-Lagrangian evaluation cost is thus $O(EN^{2d})$. Moreover, the constant is not small, given that several interpolations are required, the Lagrangian bases need to be constructed, and, for deformed geometries, Newton's method must be applied to find the computational coordinates in the reference element as a function of \mathbf{X}_j^{n-q} . This cost must be contrasted to the cost of Eulerian operator evaluations which have operation count scaling as $O(EN^{d+1})$ and relatively small constants (e.g., for the spectral element method, one can evaluate $w = \mathbf{c} \cdot \nabla \phi$ in approximately $2dEN^{d+1}$ operations).

5.5 The Operator-Integration Factor Scheme

The operator-integration factor scheme (OIFS), due to Maday, Patera, and Rønquist [?] can be viewed as a variant of the characteristics scheme (5.11) in which the values of $\tilde{\phi}^{n-q}$ are obtained without using off-grid interpolation. The basic idea is to recognize that to evaluate $\tilde{\phi}^{n-q}$, one does not need to know \mathbf{X}^{n-q} as (5.12) would seem to indicate. Rather, one can find $\tilde{\phi}^{n-q}$ by simply solving the following initial-boundary value problem:

$$\begin{aligned} \frac{\partial \tilde{\phi}}{\partial s} + \mathbf{c} \cdot \nabla \tilde{\phi} &= 0, & s \in [t^{n-q}, t^n] \\ \tilde{\phi}(\mathbf{x}, t^{n-q}) &= \phi(\mathbf{x}, t^{n-q}) & \tilde{\phi}(\mathbf{x}, t) = \phi(\mathbf{x}, t) \quad \forall \mathbf{x} \in \partial\Omega_{\mathbf{c}}. \end{aligned} \quad (5.15)$$

Here, $\partial\Omega_{\mathbf{c}}$ is defined as the subset of the boundary $\partial\Omega$ where $\mathbf{c} \cdot \mathbf{n} < 0$, that is, the portion of $\partial\Omega$ having incoming velocity characteristics. (Note that, in practice, one can assign $\tilde{\phi}(\mathbf{x}, t) = \phi(\mathbf{x}, t^{n-q})$ on $\partial\Omega_{\mathbf{c}}$. PF, verify.)

Equation (5.15) is a pure convection problem. It has the effect of propagating the initial condition (in this case, the values of ϕ at time t^{n-q}) forward, along the characteristics of the convecting field \mathbf{c} . After a time of $q\Delta t$, the values of $\tilde{\phi}$ at the gridpoints \mathbf{x}_j are precisely the desired values $\tilde{\phi}^{n-q}(\mathbf{x}_j) := \phi(\mathbf{X}_j^{n-q}, t^{n-q})$. These values can be inserted directly into (5.14). Note that *no* interpolation is required to find $\tilde{\phi}^{n-q}$ because (5.15) is solved in the Eulerian framework, using explicit time marching with a stepsize, $\Delta s \leq \Delta t$, that is sufficiently small to satisfy the *CFL* constraint. Because of the short duration of the time integration ($q\Delta t$), it is preferable to solve (5.15) using a multistage integrator such as Runge-Kutta, rather than a multistep integrator such as ABk, to avoid the start-up problems associated with lack of information prior to the first time step. Note that if \mathbf{c} is a function of t and known only at discrete times t^{n-q} , then interpolation will need to be used to approximate it at intermediate times (e.g., $t^{n-q} + \Delta s$, $t^{n-q} + 2\Delta s$, etc.). Fornberg's interpolation routine is also useful in this context.

In summary, the OIF scheme involves solving (5.15) k times, with different initial conditions and over different time intervals; summing these solutions together with weights β_q to construct the right-hand side (5.14); and finally solving the Helmholtz problem (5.13) to advance the solution to the next step.

As it stands, the OIF scheme has a complexity scaling as $O(k^2)$, for there are k integrations to be performed over time intervals ranging from Δt to $k\Delta t$ in length. For example, suppose that one wishes to use OIFS with BDF3, using a Δt -based CFL of 5. To compute $\tilde{\phi}^{n-3}$, one needs to integrate (5.15) from t^{n-3} to t^n , or a total of $3\Delta t$ in time. Using RK4, whose stability curve intersects the imaginary axis at $\lambda\Delta t \approx 2i$, and assuming, with some margin of safety, that the maximum eigenvalue of the spectral element discretization satisfies $\lambda_{\max}\Delta t \approx 2CFL$, one should set the substep size $\Delta s \approx \Delta t/5$. That is, five RK4 steps, each requiring four function evaluations, would be required for *each* Δt interval covered by the integration. The total number of function evaluations would thus be $5 \times 4 \times 3 = 60$ for the computation of $\tilde{\phi}^{n-3}$. For $\tilde{\phi}^{n-2}$, one would have 40, and for $\tilde{\phi}^{n-1}$, one would have 20.

This complexity can be reduced to $O(k)$ by recognizing that (5.15) is *linear*. Therefore superposition can be used to compute the k

solutions in a single pass. We replace (5.15) with the following system: For $q = k, k-1, \dots, 1$:

$$\begin{aligned} \frac{\partial \psi^q}{\partial s} + \mathbf{c} \cdot \nabla \psi^q &= 0, \quad s \in [t^{n-q}, t^{n-q+1}] \\ \psi(\mathbf{x}, t^{n-q}) &= \psi^{q+1}(t^{n-q}) + \beta_q \phi(\mathbf{x}, t^{n-q}) \quad \psi(\mathbf{x}, t) = \phi(\mathbf{x}, t) \quad \forall \mathbf{x} \in \partial\Omega_{\mathbf{c}}. \end{aligned} \quad (5.16)$$

With $\psi^{k+1} := 0$, the final result is

$$\psi^1(t^n) = \sum_{q=1}^k \beta_q \tilde{\phi}^{n-q}, \quad (5.17)$$

which is the desired contribution to the right-hand side of (5.14).

5.6 Extension to Navier-Stokes

The extension of the characteristics/OIF method to the Navier-Stokes equations is straightforward. Following the formalism used above for the convection-diffusion equations, we write (5.1) as

$$\begin{aligned} \frac{Du}{Dt} &= -\nabla p + \frac{1}{Re} \nabla^2 \mathbf{u} \quad \text{in } \Omega, \\ \nabla \cdot \mathbf{u} &= 0 \quad \text{in } \Omega. \end{aligned} \quad (5.18)$$

The (here nonlinear) convective term has again been expressed as a material derivative. As in the convection-diffusion problem, the objective is to replace the conditionally stable treatment of the convective term with an unconditionally stable BDF k treatment of the material derivative. (Note that large values of k are of particular interest here since one needs to compensate for accuracy loss resulting from larger values of Δt .) Using this approach, the semidiscretized Navier-Stokes equations for $k = 2$ are

$$\begin{aligned} \frac{3\mathbf{u}^n - 4\tilde{\mathbf{u}}^{n-1} + \tilde{\mathbf{u}}^{n-2}}{2\Delta t} &= -\nabla p^n + \frac{1}{Re} \nabla^2 \mathbf{u}^n \quad \text{in } \Omega, \\ \nabla \cdot \mathbf{u}^n &= 0 \quad \text{in } \Omega. \end{aligned} \quad (5.19)$$

Again, the values $\tilde{\mathbf{u}}^{n-q}$ represent the field \mathbf{u} at the foot of the characteristics. These may be computed using the usual semi-Lagrangian formulation involving off-grid interpolation, or by the OIF approach, in which one solves the convective subproblems

$$\begin{aligned} \frac{\partial \tilde{\mathbf{u}}}{\partial s} + \mathbf{u} \cdot \nabla \tilde{\mathbf{u}} &= 0, \quad s \in [t^{n-q}, t^n] \\ \tilde{\mathbf{u}}(\mathbf{x}, t^{n-q}) &= \phi(\mathbf{x}, t^{n-q}) \quad \tilde{\mathbf{u}}(\mathbf{x}, t) = \phi(\mathbf{x}, t) \quad \forall \mathbf{x} \in \partial\Omega_{\mathbf{c}}, \end{aligned} \quad (5.20)$$

where, $\partial\Omega_c$ is defined as the subset of the boundary $\partial\Omega$ where $\mathbf{u} \cdot \mathbf{n} < 0$.

Note that there should be no confusion in (5.20) between the *convected* field, $\tilde{\mathbf{u}}$, and the *convecting* field, \mathbf{u} , which is the solution to (5.19). In solving (5.20), one treats $\tilde{\mathbf{u}}$ as a passive advected vector field. Values of \mathbf{u} on the interval $[t^{n-k}, t^n]$ are computed by interpolating from the known fields $\mathbf{u}^{n-k}, \mathbf{u}^{n-k+1}, \dots, \mathbf{u}^{n-1}$, which guarantees a divergence-free convecting field in (5.20). Unlike the convection-diffusion problem, however, one does not have access to \mathbf{u}^n , so \mathbf{u} is effectively *extrapolated* on $(t^{n-1}, t]$. One has the option of iterating between (5.20) and (5.19) to get an estimate of \mathbf{u}^n , which can then be employed in the interpolation step required when solving (5.20) on a second (third, etc.) pass. Formally, this should not be required, as the extrapolation-based scheme is $O(\Delta t^k)$ accurate. However, it may be of interest in particularly challenging applications. If properly coded, (e.g., by relying on general interpolation routines such as Fornberg's), this iterative approach can be included as an option in the OIFS implementation from the outset with little overhead.

Remark on Steady State Error

We comment that one advantage of the BDF k /EXT k formulation is that there is no temporal error for problems that have evolved to a steady state solution. That is, the final result (and error) is independent of Δt . To see this, consider (5.4) under steady-state conditions, $\phi^n = \phi^{n-1} = \dots$. The lefthand side of (5.4) vanishes identically and represents the only term containing Δt . The righthand side vanishes through a balance of the diffusion and advection terms that are independent of Δt and the error will (ultimately) reflect only the spatial discretization parameters.

On the other hand, the OIFS formulation *does* retain a steady state error that is influenced by Δt . This is evident from (5.11). Under steady state conditions, we again have $\phi^n = \phi^{n-1} = \dots$, but these are not the terms that are required to balance. The lefthand side of (5.11) involves weighted differences of $\phi^n, \tilde{\phi}^{n-1}$, and $\tilde{\phi}^{n-2}$ that are divided by Δt . In fact, there should be no expectation that the lefthand side should vanish, even as $\Delta t \rightarrow 0$, because the lefthand side represents the variation of ϕ as one travels along the characteristic. This material derivative will generally be nonzero and the temporal finite difference scheme will retain an $O(\Delta t)$ dependence even under steady-state conditions. The characteristics-based OIFS scheme amounts to mixing spatial and temporal discretizations and

care must be exercised when implementing such methods, particularly in a high-order context.

Appendix A: Stability Region Evaluation

The matlab code used to generate Fig. 5.1 is presented below. The neutral stability curves are generated by considering application of a particular timestepping scheme to the homogeneous model problem

$$\frac{du}{dt} = \lambda u.$$

Assume that a uniform stepsize, Δt , is employed and that the timestepping scheme is of the form:

$$\frac{1}{\Delta t} \sum_{q=0}^{k_1} \beta_q u^{n-q} = \lambda \sum_{p=1}^{k_2} \gamma_p u^{n-p}. \quad (5.21)$$

Eq. (5.21) has solutions of the form

$$u^n = (z)^n u^0, \quad (5.22)$$

where $(.)$ is used to indicate that the (generally complex) argument z is raised to the n th power. In order for (5.22) to be neutrally stable, we require $|z| = 1$, that is, $z = e^{i\theta}$ for some value of θ between 0 and 2π . Inserting (5.22) into (5.21) and solving for $\lambda\Delta t$ yields

$$\lambda\Delta t = \frac{\sum_{q=0}^{k_1} \beta_q e^{-iq\theta}}{\sum_{p=1}^{k_2} \gamma_p e^{-ip\theta}}. \quad (5.23)$$

Evaluating (5.23) for $0 \leq \theta < 2\pi$ yields the locus associated with the neutral stability curve.

The matlab code below fills a vector `ei` with complex values on the unit circle, then creates column vectors `exp(q*ei)` for `q=1,0,-1,-2`. Linear combinations of these column vectors form the numerator and denominator of (5.23). The matlab `plot` routine is used to plot the locus of complex pairs.

Matlab Code for Stability Analysis

```

ii=sqrt(-1);
th=0:.001:2*pi; th=th';
ith=ii*th;
ei=exp(ith);
E = [ ei 1+0*ei 1./ei 1./(ei.*ei) ];

ep = 1.e-13

ab0 = [1 0.0 0.0 0.]';
ab1 = [0 1.0 0.0 0.]';
ab2 = [0 1.5 -1.5 0.]';
ab3 = [0 23./12. -16./12. 5./12.]';
bdf1 = (([ 1. -1. 0. 0.])/1.)';
bdf2 = (([ 3. -4. 1. 0.])/2.)';
bdf3 = (([11. -18. 9. -2.])/6.)';
ex0 = [0 1 0 0]';
ex1 = [0 2 -1 0]';
ex2 = [0 3 -3 1]';
yaxis = [-1.0*ii 1.0*ii]';
xaxis = [-2.0+ep*ii 2.0+ep*ii]';
du = [1. -1. 0. 0. ]';

hold off; plot (yaxis,'k-'); axis square; axis([-1.5 0.5 -1 1]); hold on;
plot (xaxis,'k-');
ab3 = (E*du)./(E*ab3); plot (ab3 ,'r-');
bdf3ex2 = (E*bdf3)./(E*ex2); plot (bdf3ex2,'k-');
print -deps ab3bdf3.eps

hold off; plot (yaxis,'k-'); axis square; axis([-1.5 0.5 -1 1]); hold on;
plot (xaxis,'k-');
ab2 = (E*du)./(E*ab2); plot (ab2 ,'r-');
bdf2ex1 = (E*bdf2)./(E*ex1); plot (bdf2ex1,'k-');
print -deps ab2bdf2.eps

```